

PATENT APPLICATION

APPARATUS AND METHOD FOR ONLINE DATA MIGRATION WITH REMOTE COPY

Inventors: **Akira Fujibayashi**
4685 Albany Circle, #118
San Jose, CA 95129
Citizenship: Japan

Akira Yamamoto
1202 Ruppell Place
Cupertino, CA 95014
Citizenship: Japan

Assignee: **Hitachi, Ltd.**
6, Kanda-Surugadai, 4-Chome
Chiyoda-ku
Tokyo 101-8010
JAPAN
Incorporation: Japan

Entity: Large

Please direct communications to:
Squire, Sanders & Dempsey, L.L.P.
600 Hansen Way
Palo Alto, CA 94304-1043
(650) 856-6500

EXPRESS MAIL LABEL NO.: EL 701 365 001 US

APPARATUS AND METHOD FOR ONLINE DATA MIGRATION WITH REMOTE COPY

This application is a continuation of U.S. Patent Application No. 09/927,167 filed August 10, 2001 by Akira Fujibayashi, et al. which is hereby incorporated by reference.

BACKGROUND OF THE INVENTION

1. Field Of The Invention

This invention relates generally to computer systems, and more particularly to systems, methods, and associated apparatus for computer data migration.

2. Description Of The Background Art

Online data migration is the process of transferring data from an old storage system to another while data stored in the old storage system remain accessible. Online data migration has become popular with computer users because it allows for storage system replacements with minimum impact on services. Online data migration, in general, is known in the art; e.g., see the paper "CARE Software Solutions, Hitachi On-line Data Migration", available from Hitachi Data Systems of Santa Clara, California, incorporated herein by reference in its entirety.

"Remote copy" is a mirroring technology that involves the copying of data from a local site to a remote site for back-up purposes. Typically, remote copy is done in real-time to provide continuous data protection, and thereby minimize data loss in the event of a disaster such as a storage system crash or unscheduled facility shutdown (e.g., due to fire, earthquake, power loss, etc.). Remote copy software and associated apparatus have been developed for backing-up the storage systems of a local mainframe computer using remote storage systems.

Computer users who have already implemented remote copy face a dilemma when replacing their storage systems during an online data migration. On the one hand,

these users typically have a need to upgrade their old storage systems in order to take advantage of the faster speed and higher capacity of newer storage systems. On the other hand, they cannot afford to disrupt their existing remote copy functions, and risk the possibility of data loss, during the online data migration process. Because currently
5 available techniques for performing online data migration require that the remote copy function be stopped for long periods of time and/or reconfigured after the migration (which also takes a long time because the remote copy has to start over from an initial copy), most users do not have a choice but to risk data loss during the migration process.

10 From the foregoing, a technique for performing online data migration with minimum impact on remote copy is highly desirable.

SUMMARY

The present invention relates to methods, systems, and apparatus for performing online data migration with remote copy. In one embodiment, a new storage system is
15 connected to an old storage system destined for replacement. Prior to starting online data migration, remote copy configuration information from the old storage system is migrated to the new storage system. This allows either the new storage system or the old storage system, depending on configuration, to maintain an operational remote copy function with another storage system during the online data migration.

20 In one embodiment, the old and new storage systems are primary storage systems located in a local site. In that case, during the online data migration, write data destined for the old primary storage system are written to both the old and new primary storage systems to maintain data integrity.

In one embodiment, the old and new storage systems are secondary storage systems located in a site remotely located from a local site.

These and other features and advantages of the present invention will be readily apparent to persons of ordinary skill in the art upon reading the entirety of this disclosure, which includes the accompanying drawings and claims.

DESCRIPTION OF THE DRAWINGS

FIG. 1 schematically illustrates an online data migration setup for replacing a primary storage system in accordance with an embodiment of the present invention.

FIG. 2 shows a flow diagram of a method for replacing a primary storage system in accordance with an embodiment of the present invention.

FIG. 3 schematically illustrates the final configuration of storage systems after an online data migration in accordance with an embodiment of the present invention.

FIG. 4 schematically illustrates an online data migration setup for replacing a primary storage system in accordance with another embodiment of the present invention.

FIG. 5 shows a flow diagram of a method for replacing a primary storage system in accordance with another embodiment of the present invention.

FIG. 6 schematically illustrates an online data migration setup for replacing a secondary storage system in accordance with an embodiment of the present invention.

FIG. 7 shows a flow diagram of a method for replacing a secondary storage system in accordance with an embodiment of the present invention.

FIG. 8 schematically illustrates an online data migration setup for replacing a primary storage system and a secondary storage system in accordance with an embodiment of the present invention.

FIG. 9 shows a flow diagram of a method for replacing a primary storage system and a secondary storage system in accordance with an embodiment of the present invention.

DETAILED DESCRIPTION

FIG. 1 schematically illustrates an online data migration setup for replacing a primary storage system with another from the same manufacturer, in accordance with an embodiment of the present invention. In the setup of FIG. 1, host processor 101, old primary storage system 102, and new primary storage system 103 are located in a local site, whereas old secondary storage system 104 is located in a remote site. As used throughout the present disclosure, the term "primary storage system" refers to a storage system located in a local site, while the term "secondary storage system" refers to a back-up storage system located in a remote site.

In FIG. 1, host processor 101 is a mainframe-type host processor. It should be understood, however, that the invention is not so limited, and is applicable to computer systems in general. Old primary storage system 102 and new primary storage system 103 are both from the same manufacturer (e.g., Hitachi Data Systems), and are, generally speaking, hardware and/or software compatible. Storage systems, in general, have a controller and one or more storage units. The controller, which directs the operation of the storage system, includes a set of channel interfaces for communicating with the host processor and other storage systems, a local processor and associated

memory for running internal processes, a cache controller and associated cache memory for performing caching operations, and a storage unit interface for communicating with internal storage units. The storage units contain media for storing data (e.g., tape, magnetic disk, optical disk). In the present embodiment, the storage units are disk units, and the storage media are either optical or magnetic disks. Storage systems are available from various manufacturers including Hitachi Data Systems.

Still referring to FIG. 1, old primary storage system 102 is to be replaced with new primary storage system 103. Prior to its replacement, old primary storage system 102 was coupled to host processor 101 via a connection 105; data on old primary storage system 102 were backed-up on old secondary storage 104 using a remote copy function over connection 106. As used throughout the present disclosure, the term "connection" generally refers to a communications link, and may be in-bound (e.g., a mainframe I/O channel) or out-bound (e.g., an Ethernet link). A connection can be moved in a variety of ways including by re-routing or replacing the physical cables making the connection, or by making logical changes (e.g., by changing the value of registers, pointers, memory locations) to redirect data to a new communications path.

An on-line data migration using the setup of FIG. 1 is now described with reference to the method of FIG. 2. Referring to action 202 of FIG. 2, replacing old primary storage system 102 with new primary storage system 103 begins by connecting old primary storage system 102 to new primary storage system 103 via a connection 108. Because new primary storage system 103 is not yet activated, action 202 can be performed without disrupting the operation of host processor 101 and the remote copy function between old primary storage system 102 and old secondary system 104. In

action 204, the remote copy function between old primary storage system 102 and old secondary storage system 104, as well as communications between old primary storage system 102 and host processor 101 are disrupted.

5 In action 206, connection 105 between host processor 101 and old primary storage system 101 is removed, and a connection 109 (see FIG. 1) between host processor 101 and new primary storage system 103 is established.

In action 208, remote copy configuration information is migrated from old primary storage system 102 to new primary storage system 103 via connection 108. Generally speaking, "migration" is the process of transferring data from one storage system to
10 another. Migration may involve copying of data as-is if both storage systems are functionally identical; otherwise, some data format conversion may be necessary. Because old primary storage system 102 and new primary storage system 103 are from the same manufacturer in the instant case, the migration process is relatively straightforward and the remote copy configuration information is migrated as-is (with
15 minor modifications to reflect the new copy-pair as explained below).

The remote copy configuration information specifies the setup for copying data from a primary storage system to a secondary storage system. It includes a copy-pair information, which identifies a storage system in the remote site designated to "mirror" a storage system in the local site. An example copy -pair information is shown in Table
20 1.

TABLE 1. EXAMPLE COPY-PAIR INFORMATION

Local Storage System
Local Logical Storage Unit

Remote Storage System
Remote Logical Storage Unit

As shown in Table 1, a copy-pair information also specifies which logical storage unit in a particular storage system in the local site is paired with a logical storage unit in a particular storage system in the remote site. As part of the migration of the remote copy configuration information, the copy-pair information is updated to reflect the new pair of storage systems. The storage systems, and the logical storage units in the storage systems, are identified by their respective addresses and other location information such as the connection between the local and remote storage systems. During remote copy, data written on a local logical storage unit are copied to the corresponding remote logical storage unit. Because data caching is employed in the storage systems, the copying may be performed right away or at some later time.

The remote copy configuration information also includes a control bitmap and some management information. The control bitmap has information showing the difference between data on the local and remote storage systems. In one embodiment, the control bitmap includes information indicating which data blocks have been updated and thus need to be copied to the remote storage system. The management information includes the status of the pair of storage systems. In one embodiment, the pair of storage systems can either be in the SUSPEND, PAIR, or COPY state. SUSPEND indicates that the remote copy process between the local and remote storage systems is suspended. The difference between the local and remote storage systems, which is reflected in the control bitmap, increases in the SUSPEND state. PAIR indicates that there is no difference between the local and remote storage

systems, thus allowing write requests to be processed immediately (synchronous mode) or periodically (asynchronous mode). COPY indicates that data are being copied from the primary storage system to the corresponding secondary storage system.

Referring to action 210 of FIG. 2, connection 106 is removed and connection 107 (see FIG. 1) is established after the migration of the remote copy configuration information from the old primary storage system 102 to new primary storage system 103. In action 212, a remote copy function between new primary storage system 103 and old secondary storage 104 over connection 107 is started.

Also in action 212, migration of the rest of the data of old primary storage system 102 begins. During migration, requests to read data locations not yet migrated from old primary storage system 102 are serviced by new primary storage system 103 by reading the requested data from old primary storage system 102 over connection 108; new primary storage system 103 then provides the requested data to host processor 101 over connection 109. Write requests are serviced by new primary storage system 103 by writing the accompanying write data to both itself and old primary storage system 102 to maintain data integrity.

In action 214, old primary storage system 102 is removed once all of its data have been migrated to new primary storage system 103. Of course, once the migration is complete, the system begins to operate normally, and special migration-related processes such as the process for writing to both the new and old primary storage systems are terminated. FIG. 3 schematically illustrates the final arrangement of host processor 101, new primary storage system 103, and old secondary storage system

104 after the online data migration has completed and the old primary storage system 102 has been accordingly removed.

5 The online data migration setup of FIG. 1 and the method of FIG. 2 allow for minimum disruption of the remote copy function. In the method of FIG. 2, remote copy is stopped in action 204, and is operational again in action 212. Because the migration of remote copy configuration information in action 208, and the movement of connections in actions 206 and 210 can be completed in a relatively short amount of time, the remote copy function is only down for a short period. Furthermore, the remote copy between the new primary storage system and the old secondary storage system 10 does not have to be reconfigured, and is thus operational at the beginning of the online data migration in action 212.

In the online data migration setup of FIG. 1 and method of FIG. 2, the host connection and the remote copy connection are re-routed to the new primary storage system prior to the beginning of the data migration. This results in a single service 15 disruption period which starts in action 204 and ends in action 210. Depending on the application, the disruption period may also be divided into two: one for re-routing the host connection, and another for re-routing the remote copy connection. FIG. 4 schematically illustrates another online data migration setup for replacing a primary storage system with another from the same manufacturer, in accordance with an 20 embodiment of the present invention. FIG. 4 is now described with reference to the method of FIG. 5. In action 502, replacing old primary storage system 102 with new primary storage system 103 begins by connecting old primary storage system 102 to new primary storage system 103 using connection 108 (see FIG. 4). Because new

primary storage system 103 is not yet activated, action 502 can be performed without disrupting the operation of host processor 101, and the remote copy function between old primary storage system 102 and old secondary system 104. In action 504, communications between old primary storage system 102 and host processor 101 are
5 disrupted. In action 506, connection 105 between host processor 101 and old primary storage system 102 is removed, and connection 109 (see FIG. 4) between host processor 101 and new primary storage system 103 is established.

In action 508, normal operations are resumed and the online data migration is started. That is, new primary storage system 103 can begin processing read and write
10 requests from host processor 101 after connection 109 is established. During this time, requests to read data locations not yet migrated from old primary storage system 102 are serviced by new primary storage system 103 by reading the requested data from old primary storage system 102 over connection 108; new primary storage system 103 then provides the requested data to host processor 101 over connection 109. Write requests
15 are serviced by new primary storage system 103 by writing the accompanying write data to both itself and old primary storage system 102. This maintains data integrity during the migration process, and allows the remote copy function between old primary storage system 102 and old secondary storage system 104 to keep on running. Note that connection 106, which is the remote copy connection between old primary storage
20 system 102 and old secondary storage system 104, has not yet been disrupted.

At some point in time, all data, including the remote copy configuration information, of old storage system 102 will be migrated to new primary storage system 103. In action 510, new primary storage system 103 then begins treating all remote

copy write data as pending data to be copied to old secondary storage system 104; this saves remote copy data during the time the remote copy function is disrupted. In action 512, the remote copy function between old primary storage system 102 and old secondary storage system 104 is disrupted. Optionally, old primary storage system 102 can be removed at this time. In action 514, the remote copy connection is re-routed by removing connection 106, and establishing connection 107 between new primary storage system 103 and old secondary storage system 104 as shown in FIG. 3. In action 516, remote copy is started between new primary storage system 103 and old secondary system 104, ending the online data migration process. As before, the remote copy does not have to be reconfigured.

The present invention may also be used in replacing an old secondary storage system with another from the same manufacturer. FIG. 6 schematically illustrates an online data migration setup for such an application, in accordance with an embodiment of the present invention. In FIG. 6, old secondary storage system 104 is to be replaced with a new secondary storage system 601, both of which are from the same manufacturer.

The setup of FIG. 6 is now described with reference to the method of FIG. 7. In action 702, new secondary storage system 601 is connected to old secondary storage system 104 via connection 602. As can be appreciated, action 702 can be performed without disrupting the remote copy function between old primary storage system 102 and old secondary storage system 104 over connection 106. In action 704, old primary storage system 102 begins treating all remote copy write data as pending data to be copied to new secondary storage system 601; this saves remote copy data during the

time the remote copy function is disrupted. In action 706, the remote copy function between old primary storage system 102 and old secondary storage system 104 is disrupted. In action 708, the remote copy configuration information of old secondary storage system 104 is migrated to new secondary storage system 601. In action 710, 5 the remote copy connection is re-routed by removing connection 106, and establishing connection 603 between old primary storage system 102 and new secondary storage system 601. In action 712, online data migration from old secondary storage system 104 to new secondary storage system 601 begins. Also in action 712, the remote copy function between old primary storage system 102 and new secondary storage system 10 601 is started. In action 714, old secondary storage system 104 is removed after the completion of the online data migration.

As can be appreciated, replacing both the old primary storage system and the old secondary storage system in the same migration process can be accomplished by combining the just described techniques for replacing individual storage systems in the 15 local and remote sites. FIG. 8 schematically illustrates an online data migration setup for replacing old primary storage system 102 with new primary storage system 103, and old secondary storage system 104 with new secondary storage system 601 in accordance with an embodiment of the present invention. In one embodiment, all the storage systems depicted in FIG. 8 are from the same manufacturer (as will be 20 described later on below, the online data migration setup of FIG. 8 may also be used in replacing a storage system with another from a different manufacturer). FIG. 8 is now described with reference to the method of FIG. 9. In action 902, new primary storage system 103 is connected to old primary storage system 102 via connection 108, and

new secondary storage system 601 is connected to old secondary storage system 104 via connection 602 (see FIG. 8). In action 904, the remote copy function between old primary storage system 102 and old secondary storage system 104, as well as communications between old primary storage system 102 and host processor 101 are
5 disrupted.

In action 906, connection 105 between host processor 101 and old primary storage system 101 is removed, and connection 109 (see FIG. 8) between host processor 101 and new primary storage system 103 is established.

In action 908, remote copy configuration information is migrated from old primary
10 storage system 102 to new primary storage system 103 via connection 108. Also in action 908, remote copy configuration information is migrated from old secondary storage system 104 to new secondary storage system 601 via connection 602.

In action 910, the new remote copy connection is established by removing connection 106, and establishing connection 107 (see FIG. 8) between new primary
15 storage system 103 and new secondary storage system 601.

In action 912, online data migration processes from old primary storage system 102 to new primary storage system 103, and from old secondary storage system 104 to new secondary storage system 601 are started. Also in action 912, remote copy between new primary storage system 103 and new secondary storage system 601 over
20 connection 107 is started. Note that because remote copy configuration information has been previously migrated to the new storage systems, remote copy is available as soon as connection 107 is established. At this point, new primary storage system 103 can begin processing read and write requests from host processor 101. Requests to read

data locations not yet migrated from old primary storage system 102 are serviced by new primary storage system 103 by reading the requested data from old primary storage system 102 over connection 108; new primary storage system 103 then provides the requested data to host processor 101 over connection 109. Write requests
5 are serviced by new primary storage system 103 by writing the accompanying write data to both itself and old primary storage system 102. This maintains data integrity during the data migration and remote copy function.

In action 914, the old storage systems can be removed as soon as their respective online data migration ends.

10 The present invention can also be used in replacing old storage systems with new storage systems from a different manufacturer. For example, the online data migration setup of FIG. 8 and the method of FIG. 9 can also be used in online data migration between storage systems from different manufacturers. In that case, old primary storage system 102 and old secondary storage system 104 will be from one
15 manufacturer, and new primary storage system 103 and new secondary storage system 601 will be from another. The setup and method will be similar to that shown in FIGS. 8 and 9, except that the migration of remote copy configuration in action 908 involves a prior action of converting the remote copy configuration of the old storage systems to the same format as that used by the new storage systems. Manufacturers typically
20 provide a system management tool for accessing remote copy configuration information, which can then be reformatted manually via a service console or automatically using a script. The reformatting of remote copy configuration information can be performed before the new storage systems are brought onsite (i.e., before action

902). Once reformatted, the remote copy configuration can then be migrated from the old storage systems or simply copied over to the new storage systems in action 908.

Methods, systems and apparatus for performing online data migration with remote copy have been disclosed. While specific embodiments have been provided, it is to be understood that these embodiments are for illustration purposes and not limiting. For example, although the provided exemplary embodiments demonstrate the replacement of a single storage system, the present invention can be used in replacing multiple storage systems. Because many additional embodiments will be apparent to persons of ordinary skill in the art reading this disclosure, the present invention is limited only by the following claims.